

ATTORNEY DOCKET 90204

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

PATENT APPLICATION

FOR:

SEU AND SEFI FAULT TOLERANT COMPUTER

INVENTOR:

DAVID R. CZAJKOWSKI

ENCINITAS, CA

Robert P. Cogan

Reg. No. 25,049

NATH & ASSOCIATES LLP

11455 El Camino Real, Suite 210

San Diego CA 02130-2008

858-792-8211

EXPRESS MAIL LABEL
ER 967094565 US

SEU AND SEFI FAULT TOLERANT COMPUTER

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This patent application claims priority of provision patent application 60/442,727, filed January 28, 2003.

BACKGROUND OF THE INVENTION

[0002] The present invention relates to fault tolerant computers and more specifically to a method and apparatus for operating in an error free manner when a microprocessor error is induced.

[0003] Three basic factors contributing to the functioning of a computer, and more specifically to a microprocessor or microprocessors included in a computer are power, performance and environment-induced radiation effects. New models or generations of computers seek to achieve higher performance at lower power levels. Additionally, in applications in which microprocessors are exposed to ionizing radiation, it is necessary to provide a mechanism for maintaining reliable operation when it is a virtual certainty that the ionizing radiation will cause processor errors. An example of applications in which sufficient levels of radiation will be encountered to cause errors in spaceborne computers. In applications in which particle or ionizing radiation is not present, errors can be caused by other fault mechanisms such as electrically induced noise pulses.

[0004] The most significant error events are Single Event Upset (SEU) and Single Event Functional Interrupt (SEFI). SEU is defined by NASA as "radiation-induced errors in microelectronic circuits caused when charged particles (usually from the radiation belts or from cosmic rays) lose energy by ionizing the medium through which they pass, leaving behind a wake of electron-hole pairs. In other words SEU is a change of state or transient induced by an energetic particle such as a cosmic ray or proton in a device. SEUs are "soft errors" in that a reset or rewriting of the device causes normal device behavior thereafter. However, the error must be accounted for when it is included in data to be acted upon. An SEFI is a condition in which an SEU in the device's control circuitry places the

device into a test mode, halt, or undefined state. The SEFI halts normal operations, and is believed to require a power reset to recover. SEU error rates in a nominal application for commercial microprocessors can range from 0.2 to 9 MeV/mg/cm². This range of rates is reflected in processor performance, depending on the processor and its environment, from a quite acceptable single upset per year to an unacceptable multiple upsets per hour.

[0005] Improved SEU performance when designing microprocessor systems commonly results in increased power consumption. However, this technique does not solve the problem of SEUs and SEFIs due to radiation or electrically induced noise pulses. One prior art approach comprises utilizing radiation hardened microprocessors which will not be susceptible to the errors induced by radiation. However, radiation hardened microprocessors are not available in state of the art versions. They have over the past ten years lagged non-hardened processors by two to three generations. For example, currently available radiation hardened microprocessors include a 0.35 micron SOI (Silicon on Insulator) microprocessor and a 0.25 micron bulk CMOS on EPI processor (Complimentary Metal Oxide Semiconductor on Epitaxial Layer). However, state of the art microprocessors utilize 0.13 and 0.10 geometries. Radiation hardened microprocessors also lag the state of the art in terms of MIPS (Million Instructions Per Second) capability.

[0006] Another known technique is TMR, triple modular redundancy, applied at the system level, also known as spatial redundancy. Three individual or discrete processors run instructions in parallel and synchronously. The outputs of the processors are sent to a comparator that utilizes voting logic. When an SEU occurs in one processor, the other two processors will still produce matching outputs. The comparator will pass the majority output. SEFI errors are treated as SEUs. However, the processor experiencing the SEFI will remain offline until reset or otherwise corrected. TMR triples the processor power requirements compared to a single processor. Synchronizing the processors is difficult, and operation must be slowed with respect to the speed achievable by a single processor.

[0007] Time redundancy has been employed at the system level to provide the advantage of redundancy as described above while permitting the use of a single processor. In this technique, the processor executes the same instruction three times, or two times, comparing results, and runs a third time when the results do not agree. The result, or a checksum indicative of the result, is stored and the three stored outputs are compared. Three matching results indicate the absence of an SEU. If there is an SEU, a voting circuit selects the correct result. When the SEU corrupts data, the time redundancy technique will operate correctly. However, if the SEU causes an instruction to be corrupted, the technique will not operate correctly. A bit instructing a wrong operation will cause the wrong operation to be performed all three times. SEUs are not detected and SEFIs are not corrected. An improved form of time redundancy was developed by the Stanford Advanced Research and Global Observations Satellite Project (ARGOS). This technique is described in Oh, N., P.P. Shirvani and E.J. McCluskey, "Error Detection by Duplicated Instructions In Super-scalar Processors," *IEEE Transactions on Reliability*, Vo. 49, No. 7, Sep. 2001, pp. 273-284. Many errors were corrected, but still others were not.

[0008] Another prior art alternative is to build a processor using commercial, non-radiation hardened integrated circuit process and apply known RHBD (radiation hardness by design) techniques to improve radiation hardness. Once again, as in the case of radiation hardened processors, die area is increased and operating speed are compromised. Also, while commercial switching logic utilizes simple flip-flops, RHBD logic requires latches built out of many flip-flops and further logic such as inverters. Performance comparable to commercial processors which are not radiation hardened is not provided.

[0009] Examples of an improved radiation hardened system and a time redundant system are respectively disclosed in my copending patent applications Ser. No. 10/435,626 filed May 6, 2003 entitled *Fault Tolerant Computer* and Ser. No. 10/656,720 (with a coinventor) filed September 8, 2003 entitled *Functional Interrupt Mitigation for Fault Tolerant Computer*, the disclosures of which are incorporated by reference herein. It is desirable to provide a system in which a

minimal amount of radiation hardening need be done. It is desirable to provide a system in which a time redundant system is also made space redundant, but in an efficient, reliable manner. For example, it is desired to avoid the problem of synchronizing a plurality of processors.

[0010] There is little patent literature on SEFIs. Many testing efforts with microprocessors do not report SEFIs, or "hangs." It is probable that all microprocessors will exhibit SEFIs whether they have been previously observed or not. This will include both commercial and radiation hardened devices. SEUs may take place in any transistor within a complex microprocessor. When the upset occurs in a memory location, whether a register or memory site, this can be measured and corrected. However, when the upset occurs in more subtle ways, the processor may be placed in a state from which it is not recoverable. An example is the case of an induced error in combinatorial logic or in state-machine transistors. It may be initially impossible to observe an error condition within the processor. However, the error may propagate within combinatorial logic. Other unrecoverable faults could include illegal branching, upset induced exceptions, upsets in the program counter or other unobservable faults. Work by such researchers as Dr. James W. Howard of Jackson and Tull Chartered Engineers of Washington, D.C. has demonstrated that SEFIs will occur in Pentium®, PowerPC and other processors. It is highly probable that all microprocessors will exhibit SEFIs whether they have been previously observed or not. It is therefor highly desirable to provide a way of detecting SEFIs so they may be responded to and also providing a way of responding to them.

SUMMARY OF THE INVENTION

[0011] Briefly stated, in accordance with the present invention, a method and apparatus are provided utilizing time redundancy combined with spatial redundancy in which benefits of modular redundancy are provided by in which the addition of components is minimized and in which benefits of time redundancy are provided with a minimum increase in operational complexity and in which errors not resolved by prior art time redundancy techniques are detected. SEUs are responded to. Additionally, the occurrence of SEFIs is accounted for.

[0012] A non-hardened processor is made fault tolerant to SEUs and SEFIs. A processor is provided utilizing time redundancy combined with spatial redundancy, which is also referred by applicant's trademarks time-triple modular redundancy and TTMR, using a single processor to detect and respond to SEUs. External comparison circuitry is provided in a radiation hardened module to provide "TTMR" redundancy to protect for SEU errors on input output buses. Additionally, a hardened SEFI circuit is provided to periodically send a signal to the process which, in the case of a processor not in the SEFI state, initiates production by the processor of a "correct" response. If the correct response is not received within a particular time window, the SEFI circuit initiates progressively severe actions until a reset is achieved.

[0013] Other aspects of the invention are further described below. This summary is neither exhaustive nor determinative of the scope of the present invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] The invention may be further understood by reference to the following description taken in connection with the following drawings.

[0015] Of the drawings:

[0016] Figure 1 is a block diagram of a computer constructed in accordance with the present invention including an SEU detection circuit and SEU recovery circuit;

[0017]Figure 2 is a timing diagram useful in understanding the operation of the embodiment of Figure 1;

[0018]Figure 3 is a flow diagram illustrating the operation of the SEU recovery circuit and the programmed media commanding the operation;

[0019]Figure 4 is a flow diagram illustrating an alternative operation of the SEU recovery circuit and the programmed media commanding the operation; and

[0020]Figure 5 is a flow diagram illustrating the operation of the SEFI monitoring circuit of Figure 1.

DETAILED DESCRIPTION

[0021] Figure 1 is a block diagrammatic illustration of a processor 1 communicating via a bus 3 to peripheral devices 5. The processor 1 could, for example, be included in a satellite. The peripheral devices 5 may include a communication device 7 and sensors 8. Any number of well-known input and output devices may interact with the processor 1. The term processor is used here to denote a device which functions as a computer, e.g. a Pentium microprocessor chip, and does not describe only a subcomponent such as a discrete arithmetic unit. The processor 1 will in contemplated embodiments comprise a silicon chip, but may comprise any processor subject to the Single Event Upset (SEU) and Single Event Functional Interrupt (SEFI) phenomena, whether due to radiation or noise. It should be noted that a computer to be used in accordance with the present invention need not have the particular architecture as illustrated here. There are many well-known architectures providing the operation described here. Also, since microprocessor chips have many, many subsystems, it is common that representations of identical chips may take many different forms. Commercially available chips have detailed data sheets describing units available in the chips to perform various functions. In one preferred embodiment, the processor 1 is an Equator BSP-15 processor from Equator Technologies, Inc. of Campbell, California.

[0022] The bus 3 may be interfaced to the peripheral devices 5 by a universal asynchronous receiver/transmitter (UART) 10. The processor 1 also uses a peripheral component interconnect (PCI) 12 to decouple a central processing unit 14, also coupled to the bus 3, from the relatively slow peripheral devices 5. Components of the processor 1 are coupled to communicate via the bus 3. The processor 1 comprises a main memory 18 which is a synchronous dynamic random access memory (SDRAM) 18 coupled to the bus 3. In other embodiments, other forms of dynamic storage could be utilized. The SDRAM 18 is controlled by an SDRAM controller 20. An instruction control unit 28 coupled to the bus 3 coordinates execution of program instructions. In the present embodiment, arithmetic operations are performed by an arithmetic logic unit 30.

In the BSP 15 processor, the arithmetic logic unit 30 comprises first and second units 31 and 32. A clock control 36 and memory cache 38 are also coupled to the bus 3. An SEFI control circuit 40, discussed further below, is coupled to the bus 3. SEFI circuit 40 is external to processor 1.

[0023] In the “time-triple modular redundancy” (TTMR) technique, a calculation is performed at times t_0 , t_1 , and t_2 , each time corresponding to a successive cycle of the bus 3. The results are polled for “two out of three” matching to assure a correct result. The present invention examines both memory and bus data transfers by adding an external hardware compare operation in the path of data being processed. The additional hardware should be radiation hardened. By simplifying the technique, the additional hardware, and thus the expense in its implementation is minimized. In the present invention, the computation is performed twice. The first computation is the original computation, and the second computation is referred to as a mirror calculation. If a match is obtained when the successive results produced at times t_0 and t_1 are compared, then two matching results are known to exist. It is, therefore, unnecessary to perform the third computation using the value produced at time t_2 . Since, in a nominal application, SEUs occur only about 1% of the time, it is not necessary to perform the third calculation 99% of the time.

[0024] In an SEU detection circuit 48, a comparison of first and second signals is made by a comparator 50. As used herein, discrete logic primarily refers to a “hardware” rather than “software” implementation. While logic elements in Figure 1 are illustrated as discrete logic elements, they do not need to be discrete components. The logic circuitry of Figure 1 could be embodied in a larger chip either as separately identifiable components or embodied within an integrated circuit, e.g. a field programmable gate array (FPGA). A first input is provided to the comparator 50 from a delay line 52. A second input to the comparator 50 is coupled from the SDRAM 18. The comparator provides an output to the bus 3 having a first state indicative of a match between the two inputs or a second state indicate of a non-match. The second state is referred to as an SEU error flag.

[0025] The SEU error flag initiates operation of the SEU recovery circuit 60. A first comparator 62 compares the outputs calculated at times t_1 and t_2 . A second comparator compares 64 the results produced at times t_0 and t_2 . Error flag logic circuit 66 receives the outputs of the comparators 62 and 64 to provide an output of the first state if either of the comparators 62 and 54 indicate a match. If there is not a match at either comparator 62 or 64, the error flag logic circuit produces an error signal to prohibit use of an incorrect calculation.

[0026] Operation is described with respect to Figure 2, which is a timing diagram. In Figure 2, the abscissa is time, divided into cycles of the bus 3, and the ordinate is amplitude on an arbitrary scale indicative of logical zeros or ones. Figure 2a illustrates the signal to the first input of the comparator 50, Figure 2b illustrates the second input to the comparator 50 and Figure 2c illustrates the output of the comparator 50. At time t_0 , an input indicative of a first result is supplied to the delay line 52 from the SDRAM 18 under the control of the SDRAM control 20. At time t_1 , an input indicative of a second result is supplied to the second input of the comparator 50 and also to the input of the delay line 52. By time t_1 , the first result has propagated to the first input of the comparator 50. Consequently, the comparator 50 compares the first and second results produced by the processor 1. If the inputs to the comparator agree, an output of the first state is provided by the comparator 50. This output is interpreted by the SDRAM control 20 so that the value produced by the calculation under consideration. The value is released for further processing in accordance with the programmed instructions. The input, delay and comparison process is not repeated. If the inputs do not agree, as illustrated in the example of Figure 2, then the comparator 50 produces the SEU error flag. The SEU error flag is used to call operation of the SEU recovery circuit 60.

[0027] Operation of the SEU recovery circuit is illustrated in Figure 3, which is a flow diagram. In the situation in which an SEU occurs, at block 100, the SEU error flag is produced by the comparator 50 and supplied to the processor 3 to call the operation of Figure 3. The SEU may also be detected in the absence of a value to be compared as well as in the case of a mismatch. Absence of a

signal in the present example is the failure of an input to the comparator to occur prior to a timeout, which will not exceed on bus 3 cycle. At block 102, the inputs data to the comparator 50 are each written to a storage location. The original and mirror outputs produced at t_0 and t_1 are respectively referred to as C and C'. The processor 1 is commanded at block 104 to produce two further successive outputs at successive cycles of the bus 3. The results of the initial calculation and the mirror calculation are stored as D and D' at block 106. In the nominal environment for the present invention, if there has been an SEU in the cycle in which C and C' were produced, the probabilities are such that there should not be an SEU in the calculation of D and D'. C is compared to D at block 110. C' is compared to D' at block 112. At block 110, if C matches D, the value of C is treated as "true," and the value of C is sent to bus 3 to be utilized as a valid result. At block 112, if C' matches D', the value of C' is treated as "true," and the value of C is sent to bus 3 to be utilized as a valid result.

[0028] Alternatively, the method of Figure 4 may be used to respond to an SEU error flag. At block 150, a command is issued to store instructions a each instruction is determined to be error free. At block 152, the SEU error flag is generated. At block 154, the stored commands are examined to determine the last instruction having an error free status. At block 156, the instruction control unit 28 is "decremented" to return to the last error free operation, and at block 158, the instruction stream is resumed and discarded operations are repeated.

[0029] It is also desirable to detect SEFIs. These are faults from which the processor 1 does not recover. The SEFI circuit 40 (Figure 1) is a radiation hardened circuit to monitor status of the processor 3 and reset it. As indicated in Figure 5, at block 200, the SEFI circuit 40 provides a test signal to the processor 1. The period of the test signal production may be relatively long. The test signal 1 requires processing by the CPU 14, as indicated at block 202. If the processor 1 is not in the SEFI mode, it will respond by producing a "correct" answer as indicated at block 204. The SEFI circuit 40 must receive the correct answer before a preselected time-out, such as one or a preselected number of cycles of the bus 3. As indicated at block 206, if the correct answer is received, operation

returns to block 200 to be resumed at the beginning of a next test signal period. If not, operation proceeds to block 210, where a corrective action routine is called. A first corrective action is initiated at block 212. This action is toggling of the interrupt of the CPU 14.

[0030] At block 214, operation is tested. If the processor 1 is returned to a known, operative state, the operation ceases until the next test signal. If not, operation proceeds to block 216, which is a software reboot with a flag set to signify an SEFI event. At block 218, operation is tested. If the processor 1 is returned to a known, operative state, the operation ceases until the next test signal. Also, the SEFI circuit 40 may produce a "return from SEFI" flag. In not, operation proceeds to block 220. The corrective action at block 220 is a hardware reset utilizing the "reset" input of the CPU 14. At block 222, operation is tested. If the processor 1 is returned to a known, operative state, the operation ceases until the next test signal. If not, operation proceeds to block 224 at which the CPU 14 is run through a power cycle. At block 226, operation is tested. If the processor 1 is returned to a known, operative state, the operation ceases until the next test signal. If not, operation proceeds to block 228. At block 228, the processor 1 is powered down and then restarted.

[0031] Each correction will attempt to return the CPU from SEFI by operating special software routines to self-test of roll back operation to return the hardware to a known state. The SEFI circuit 40 can be implemented by triple modular redundant FPGAs or it can be radiation hardened application specific integrated circuit (ASIC). Since the digital logic needed for the SEFI circuit 40 is estimated to be 6,000 gates, it can be implemented on a relatively small silicon chip at reasonable cost. Recovery capabilities are embedded in software routines, such as the ability to store selected data variables in memory for later recovery. Additional recovery capabilities are embedded in software routines such as the ability to store selected data variables in memory for later recovery in response to the "return from SEFI" flag.

[0032] Software embodying the above operation may be made available to users with standard software tools and languages. The most common engineering

language is C/C++. This language is supported by the Equator BSP-15 of the preferred embodiment and many widely used processors. A precompiler will duplicate computation code to produce mirror code to perform time redundant operations. The code produced for the present invention can be implemented in a real time operating system (RTOS). A preferred real time operating system is OSE_{CK} from Enea Embedded Technology of San Diego, California

[0033] The techniques of the present invention can be applied to the design of a new very long instruction word (VLIW) processor to achieve a greatly improved SEU and SEFI error rate using either hardware or software implementations. Advantageously, a microprocessor integrated circuit (IC or chip) may be designed from commercially available VLIW cores. Combined time and special redundancy and RHBD logic to a microprocessor with attention to SEU tolerance and performance will allow for significant advances in SEU hardened computing. The combined time and special redundancy can be adapted for both memory and bus data transfers by adding a hardware compare in SEU hardened logic in the data path along with the proper sequencing of data transfer and design of an SEU interrupt routine. The above teachings will enable those skilled in the art to take many departures from the specific examples above to produce systems in accordance with the present invention.